Intermediate release of the federated HPC, Cloud and storage infrastructure for EBRAINS
(D6.2 – SGA3)

Figure 1: Federated HPC, Cloud and storage infrastructure

The federated HPC, Cloud and storage infrastructure for EBRAINS is developed, deployed and operated by five European supercomputing centres, which also contribute to the Fenix infrastructure, namely BSC, CEA, CINECA, ETHZ-CSCS and JUELICH-JSC.

Project Number:	945539	Project Title:	HBP SGA3

Document Title:	Intermediate release of the federated HPC, Cloud and storage infrastructure for EBRAINS
Document Filename:	D6.2 (D58) SGA3 M21 ACCEPTED 220520.docx
Deliverable Number:	SGA3 D6.2 (D58)
Deliverable Type:	Report
Dissemination Level:	PU = Public
Planned Delivery Date:	SGA3 M21 / 31 Dec 2021
Actual Delivery Date:	SGA3 M21 / 17 Dec 2021; accepted 20 May 2022
Author(s):	Maren FRINGS, JUELICH (P20) Anna LÜHRS, JUELICH (P20)
Compiled by:	Maren FRINGS, JUELICH (P20) Anna LÜHRS, JUELICH (P20)
Contributor(s):	Cristiano PADRIN, CINECA (P13), contribution to Sections 3.1.1 and 3.1.2 Hamid FARD, TUDA (P85), contribution to Section 3.1.3 Felix WOLF, TUDA (P85), contribution to Section 3.1.3 Taylan ÖZDEN, TUDA (P85), contribution to Section 3.1.3 Björn HAGEMEIER, JUELICH (P20), contributed to Section 3.1.4 Alex UPTON, ETHZ (P18), contributed to Sections 3.1.4 and 3.2.1 Javier BARTOLOMÉ, BSC (P5), contributed to Section 3.1.5 Benedikt VON ST VIETH, JUELICH (P20), contributed to Section 3.1.6 Gilles WIBER, CEA (P11), contributed to Section 3.2.2 Wouter KLIJN, JUELICH (P20), contribution to Section 3.2.3 Claudia BACHMANN, JUELICH (P20), contribution to Section 3.2.3 Marissa DIAZ PIER, JUELICH (P20), contribution to Section 3.2.3 ... and many additional, indirect contributors from the WP6 Tasks developing and operating the services described in this document.
WP QC Review:	Boris ORTH, JUELICH (P20)
WP Leader / Deputy Leader Sign Off:	Lena ODEN, JUELICH (P20)
T7.4 QC Review:	Martin TELEFONT (tech review); Annemieke MICHELS (editorial review), EBRAINS (P1)
Description in GA:	Intermediate release of the improved and enhanced federated HPC, Cloud, storage and network infrastructure and related services for EBRAINS
Abstract:	This document summarises the intermediate release of the improved and enhanced federated HPC, Cloud, storage and network infrastructure and related services for EBRAINS, which are developed in HBP Work Package 6. It particularly focuses on the

	developments pertaining to the central data location and transfer services, a co-scheduling Slurm plugin, federated identity and access management, a monitoring service, and how to enable interactive computing. Moreover, recent activities in the areas of infrastructure operations, security assessments, use case and workflow management, and infrastructure training and webinars are briefly presented. This document concludes with a brief outlook into planned future activities.
Keywords:	Infrastructure, Cloud, High-Performance Computing, HPC, storage, network, security, monitoring, co-scheduling, Slurm, data services, interactive computing, Authentication and Authorisation Infrastructure, AAI, infrastructure operations, use case management, user training
Target Users/Readers:	Developers of platform services, infrastructure users, computer scientists, HPC community

Table of Contents

Table of Tables

Table of Figures

1. Introduction

This Deliverable summarises the intermediate release of the improved and enhanced federated HPC, Cloud, storage and network infrastructure and related services for EBRAINS that are developed and operated by the EBRAINS Computing Services Work Package (WP6). An updated version of this document will be submitted as Deliverable D6.4 “Final release of the federated HPC, Cloud and storage infrastructure for EBRAINS” at the end of the current Human Brain Project (HBP) phase in 2023.

EBRAINS Computing Services develops, deploys, integrates and operates a variety of basic IT services within the distributed HBP and EBRAINS infrastructure, including in particular the federated HPC, Cloud, storage and network infrastructure, building on resources and services of the Fenix infrastructure that are made available through the ICEI project (see Section 2 for more details about the relation to Fenix). Together with Fenix, EBRAINS Computing Services provide the computing and storage backbone of the EBRAINS infrastructure.

WP6 is also in charge of the Neuromorphic Computing systems and services. Please note that the Neuromorphic Computing systems and services are not part of this Deliverable but will be described in Deliverable D6.3 “NMC in EBRAINS with TRL8 software”, which will be submitted in 2023. The same applies to the Collaboratory, a workspace in the Cloud for research, development, documentation and collaboration, which is not part of the current Deliverable but will be described in Deliverable D6.5 “Final release of the Collaboratory” that will also be published in 2023.

This Deliverable particularly focuses on recent developments with regards to the central data location service (see Section 3.1.1), data transfer services (see Section 3.1.2), a co-scheduling Slurm plugin (see Section 3.1.3), federated identity and access management (see Section 3.1.4), a production-level monitoring service (see Section 3.1.5), and how to enable interactive computing (see Section 3.1.6). Moreover, recent activities in the areas of infrastructure operations (see Section 3.2.1), security assessments (see Section 3.2.2), use case and workflow management (see Section 3.2.3), and infrastructure training and webinars (see Section 3.2.4) are briefly described. It concludes with an outlook into planned future developments and activities (see Section 4).

This document aims at giving an overview of recent developments and activities and refers, where available, to more detailed information about the respective services and developments (e.g. to so-called “Output Documents”, websites or repositories). In includes short status updates for all Outputs (OPs), which were scheduled for the first 21 months of the current HBP phase.

2. EBRAINS Computing Services and the relation to the Fenix infrastructure

The EBRAINS Computing Services Work Package (WP6) develops, deploys, integrates and operates a variety of basic IT services within the distributed HBP and EBRAINS infrastructure, including

- the HPC/Cloud computing and storage services of the Fenix infrastructure[1], which are made available through the ICEI project;
- the neuromorphic computing services SpiNNaker and BrainScaleS;
- the HBP Collaboratory; and
- the EBRAINS accounts system ("identity provider", i.e. single sign on for EBRAINS services).

This document focuses on the first point, i.e. advances of the federated HPC, Cloud, storage and network infrastructure services. Developments of the neuromorphic systems and related services will be described in Deliverable D6.3 "NMC in EBRAINS with TRL8 software" and the Collaboratory and EBRAINS accounts system in D6.5 "Final release of the Collaboratory".

The systems and services operated by WP6 enable the platform services layer and individual vertical solutions to integrate different EBRAINS services within complex workflows. The services of this infrastructure layer (see Section 3) thus serve within the HBP and EBRAINS as a basis for the EBRAINS Data Services[2] and EBRAINS Modelling Services[3] and can, of course, also be used directly by internal and external end users. The work is mainly focused on the operation of the infrastructure services to ensure a high quality of service that can be achieved, e.g. by providing robust operational environments and establishing mechanisms that allow for timely identification of problems.

Another important focus is the adaptation of the generic Fenix infrastructure services to the specific needs of EBRAINS, e.g. regarding Authentication and Authorization Infrastructure (AAI) integration and the Fenix User and Resource Management Service (FURMS). Within the Fenix infrastructure, six European supercomputing centres, namely BSC[4] (Spain), CEA[5] (France), CINECA[6] (Italy), CSC[7] (Finland), ETHZ-CSCS[8] (Switzerland) and JUELICH-JSC[9] (Germany), have agreed to align their services. The distinguishing characteristic of this e-infrastructure is that data repositories and scalable and interactive supercomputing systems are in close proximity and well-integrated. An initial version of this infrastructure is currently being realised by BSC, CEA, CINECA, ETHZ-CSCS and JUELICH-JSC through the ICEI project (Interactive Computing E-Infrastructure), which is part of the European Human Brain Project[10] (HBP). Figure 2 displays the relation between site-local and federated infrastructure services, EBRAINS platform services like data and modelling services, and the user communities.

While the ICEI project focuses on the procurement of hardware and developments of generic (i.e. not community-specific) solutions such as FURMS, EBRAINS Computing Services operates these services and links them to the EBRAINS infrastructure by developing, enhancing and operating additional services as required by the other parts of the HBP. This document describes the recent advances in the development areas and additional important activities.

---

[1] https://fenix-ri.eu/
[2] https://ebrains.eu/services#category0
[3] https://ebrains.eu/services#category2
[4] https://www.bsc.es/
[5] http://www-hpc.cea.fr/index-en.htm
[6] https://www.cineca.it/en
[7] https://www.csc.fi/
[8] http://www.cscs.ch/
[9] https://www.fz-juelich.de/ias/jsc
[10] https://www.humanbrainproject.eu

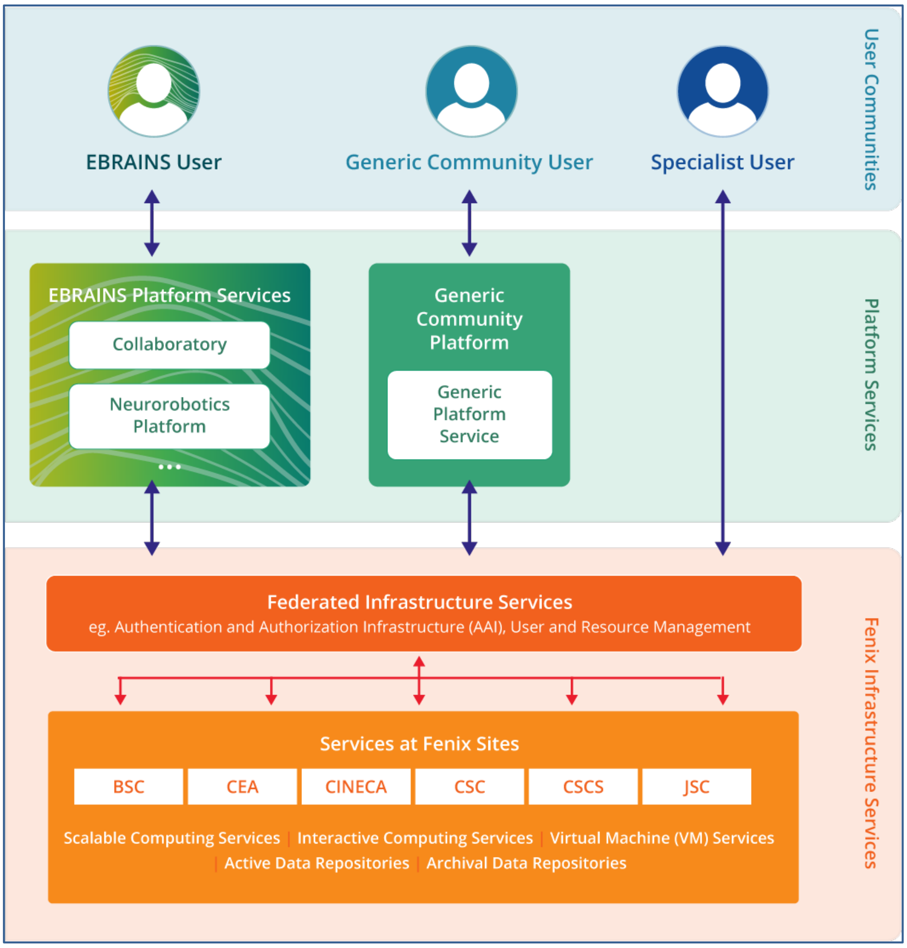

Figure 2: Relation between the Fenix infrastructure services, platform services and user communities

The Fenix sites offer various types of local infrastructure services, on top of which federated infrastructure services are deployed, which are operated by EBRAINS Computing Services. EBRAINS Platform Services run on the infrastructure services and are used by EBRAINS users. This concept can be generalised for other communities as well. Specialist users can also directly use the federated and site-local infrastructure services.

While the design of the Fenix infrastructure was to a large extent based on use cases of the HBP, collected and analysed in the ICEI project, the site-local and federated infrastructure services (i.e. all services displayed in orange) are designed to be community independent, so that they can serve multiple communities. These communities, with HBP and EBRAINS being the first large user community, then develop community-specific platform-level services, which are operated on top of the generic infrastructure services. EBRAINS Computing Services supports the development of interfaces between the generic infrastructure layer and other EBRAINS platform services with the services and activities described in this document.

To date, EBRAINS represents the only large user community, which has developed and deployed their own platform services, but several projects from different communities also successfully applied for ICEI resource allocations via the PRACE-ICEI calls[11]. Moreover, the ICEI project has established contacts with potentially interested communities. The terms “Generic Community User” and “Generic Community Platform” indicate that the Fenix infrastructure is by design not limited to the specific requirements of EBRAINS but can also serve other research communities.

[11] https://fenix-ri.eu/access (see “European Calls for Scientists of all Research Domains”)

3. New services and recent activities

This section describes all new and updated services of the federated HPC, Cloud, storage and network infrastructure, which were developed by EBRAINS Computing Services in the first 21 months of the current HBP phase, SGA3. Most of these services build on activities in previous phases of the HBP and/or the ICEI project. In addition, Section 3.2 provides an overview of the most important recent and ongoing activities in relation to the federated infrastructure.

3.1 New and updated EBRAINS Computing Services

3.1.1 Central Data Location Service

Task T6.3 will improve the data location and transfer service prototypes developed during the previous HBP phase to deliver robust and sustainable services, and to improve access to computational resources and data repositories. These services, which will connect to Fenix data repositories, will allow users to locate and retrieve data objects stored across different ICEI Archival Data Repositories (ARD) and to asynchronously transfer data between them. The transfer services are described in Section 3.1.2.

The Central Data Location service allows users to find, download, upload and share data through the usage of metadata. The service will be in a first step closely integrated with the Knowledge Graph[12] (KG). Large data sets, that are needed for or generated by EBRAINS workflows, can and should be stored in the Archival Data Repositories. The Data Location service then provides the link between the storage of the data and the access via the data in the Knowledge Graph. By using the Data Location service, the physical location of the data is automatically identified and, in the next step, transported via the Data Transfer service (see Section 3.1.2) or using the Data Mover (which will be delivered by the ICEI project to allow the site-local data transfer between object storages and close file systems) to the location where they are processed in workflow. The Data Location service therefore is the interface between the Data Services in Fenix and the Knowledge Graph. This allows the definition of workflows independent of the actual data location. In case the integration will not solve the addressed requirements, the services will be implemented in a different environment that will be developed with the same technology as the KG.

The development builds on work in the previous HBP phase, SGA2, summarised in Deliverable SGA2-D7.6.2 "SP7 High-Performance Analytics and Computing Platform – Results for SGA2 Year 2"[13].

The initial configuration of the service is almost finalised. The attributes will be defined to provide information about authors of data, version of data, scientific areas and, obviously, where data are stored. A first list of attributes is undergoing an internal evaluation process and will be approved before the end of November 2021.

The capability to store metadata of contents and data objects stored in the ARD has already been added in the past to the last official DLS JSON-LD schema. This feature has also been published and added to the HBP KG production service and is thus ready to be used. The plan is to have the future HBP Service Orchestrator (developed in HBP Work Package WP5 "EBRAINS Modelling Services") exploit such functionality, as well as the Fenix Data Mover APIs, to implement the entire workflow for the EBRAINS end users.

Progress summary and next steps

- The configuration of the service is almost finished.
- A private space in the KG should be made available to the service developers in December 2021.

[12] https://ebrains.eu/services#category0
[13] https://www.humanbrainproject.eu/en/about/governance/deliverables/sga2-phase

- The definition and evaluation of attributes and criteria should be finished before the end of November 2021.
- Work on the documentation of this service has started and should be finalised soon.
- Mid-term, an important step will be the development of advanced functionality such as file localisation across different ICEI/Fenix Archival Data Repositories.

3.1.2 Central Data Transfer Services

The Central Data Transfer Services will allow the user to transfer data between Active Data Repositories (ACD), hosted at different ICEI sites, and between Archival Data Repositories (ARD), as well hosted at different ICEI sites. The user will be able to use these services via the command line and other interfaces, to set end points and provide the token bearer to allow mapping of credentials. Task T6.3 is focusing on the development of the data transfer service between ARD, while data transfer between ACD has been deployed during the previous HBP phase (see Deliverable SGA2-D7.6.2[13] for details). This section focuses on the newly developed transfer between ARD.

A detailed description of guiding principles, use cases considered for the development, identified available technologies and conclusions and decisions drawn from these are summarised in "SGA3 - Data Transfer in Task 6.3"[14] (report for OP6.8 "Production level (TRL8) version of the Central Data Transfer service"), which will be updated in the future with additional annexes to incorporate future changes and results.

Task T6.3 has decided on the following guiding principles:

1) Allow data transfer of huge data sets between Archival Data Repositories deployed at different ICEI sites in a reliable and scalable way.
2) Foster the integration of the data transfer service within the Fenix AAI.
3) Support basic authentication mechanisms and standard protocols as OAuth2/OIDC and SAML.
4) Provide common solutions by identifying and adopting common environment settings and infrastructure configurations across sites.
5) Support the integration with different services for data management, deployed at different sites.
6) Support the integration with Cloud and Open Stack services.
7) Foster the (re-)usage of already existing technical solutions, rather than implementing new ones.

Three major use cases are considered, namely:

1) Data transfer via Command Line Interfaces (CLI): A single data object stored in an ARD at federated site A needs to be transferred asynchronously to an ARD at federated site B using the data transfer (DT) service available at site A using the CLI.
2) Bulk data transfer via CLI: A set of data objects stored in an ARD at site A needs to be transferred asynchronously to an ARD at site B using the DT service available at site A using the CLI.
3) Data transfer as part of a coordinated workflow across ICEI sites: A user has multiple allocations at all ICEI sites. This user would like to have a complex workflow that will include two or more sites. To implement his workflow, the user will do initial computations at site A and move the data to site B, possibly without any human interaction (i.e. as dependent job via the scheduler). At this point the user will be able to perform a different computation or visualisation at site B.

Several technology options have been identified, which are summarised in Table 1. More details about every technology can be found in the report for OP6.8[14] (see above).

[14] https://drive.ebrains.eu/f/6f12fc4e749543839c2c/?dl=1

Table 1: List of available technologies for data transfer

The term "Scalable" refers to the capability of a service to support an increasing number of requests, considering the total number of expected HBP users (this also includes the provided hardware resources). "Sustainable" refers to the possibility to maintain the service after the end of the current HBP phase SGA3. "Already deployed" indicates if a version of the service is already installed at all federated ICEI sites.

Technology	Scalable	Sustainable	OAuth2 supported	Already deployed
UNICORE UFTP[15]	Yes	Yes	Yes	Yes
GridFTP[16]	Yes	Maintenance after SGA3 only planned at some sites	No	Yes
Web FTS / FTS3[17]	Yes	Yes	Yes	Partially yes
XRootD[18]	Yes	Yes	No	No
OpenStack SWIFT[19]	Yes	Yes	Yes	Yes
iRODS + B2SAFE + B2STAGE[20]	Yes	Maintenance after SGA3 still be decided by some sites	Yes	Partially yes

The Data Mover service will allow the site-local data transfer between object storages and close file systems and will be delivered by the ICEI project. The Service Orchestrator is developed in HBP Work Package WP5 "EBRAINS Modelling Services".

A solution based on FTS3 to transfer data between two different object storages hosted by two different sites has already been developed and is being tested. More details about this service and about the FTS3 configuration and benchmarks results can be found in Annexes 1 and 2 of the OP6.8 report.[21]

Progress summary and next steps

- A technical solution has been identified, installed and tested at two sites.
- The best practice user guide and installation instructions have been partially deployed.
- Installation and tests will be completed at all ICEI sites in January 2022.
- Integration with AAI is in progress, and the process should be finished in the first half of February 2022.
- The first release of the service is expected in March 2022.

3.1.3 Co-scheduling Slurm plugin

New fast storage technologies, such as flash-based non-volatile memory are becoming ubiquitous in HPC systems with multiple orders of magnitude higher I/O bandwidth than traditional back-end storage systems. They can be used to heavily speed-up I/O operations, an essential prerequisite for data-intensive exascale computing capabilities. However, since the overall capacity of the fast storage available in such a multi-tiered storage system is limited, an individual job may not always benefit if access to fast storage implies longer waiting time in the queue. This is obvious if fast storage is shared across the system, such as remote shared burst buffers. The decision to use a fast

[15] https://www.unicore.eu/documentation/
[16] https://gridcf.org/gct-docs/6.2/appendices/key/index.html
[17] https://fts.web.cern.ch/
[18] https://xrootd.slac.stanford.edu/
[19] https://wiki.openstack.org/wiki/Swift
[20] https://www.eudat.eu/b2stage
[21] https://drive.ebrains.eu/f/6f12fc4e749543839c2c/?dl=1

storage tier should be supported by the batch scheduler, which can estimate when the amount of fast storage a job desires will become available.

The co-scheduling Slurm plug-in is intended to manage the allocation of user jobs to computation and storage resources in a multi-tiered storage system, according to the evaluation of I/O waiting and turnaround times in the job queue. The user specifies the required storage capacity and the amount of intermediate data read and written by the application to allow the plugin to estimate the I/O waiting and turnaround times. Based on the estimated turnaround time, the plugin decides on which storage tier the job should be allocated.

Progress summary and next steps

- A scheduling mechanism that efficiently schedules compute nodes together with capacity-bound storage in a multi-tiered storage cluster has been implemented (see publication P1949[22]).
- The latest release of the co-allocation Slurm plug-in is available in the Human Brain Project GitHub Repository[23].
- Preliminary tests to configure a temporary private parallel and fast file system were performed in a full-production HPC system located at CINECA (MARCONI100).
- Provided that, by mid of January 2022 (as proposed by CINECA), CINECA grants TUDA access to the GALILEO100 cluster, configuration and tests of the plugin on a remote-shared burst buffer are expected in the first quarter of 2022.

3.1.4 Federated Identity and Access Management service compliant with Fenix AAI for EBRAINS workflows

The ICEI architecture foresees a central system to manage membership and roles of users within certain stakeholder groups. The system is also expected to allow for resource allocation and accounting within these groups. Additional aspects, such as the management of policy documents or central management of SSH keys must also be covered. What distinguishes this e-infrastructure from prior approaches is the integration of authorisation and resource management in a central system that we call Fenix User and Resource Management Service (FURMS).

The development service for FURMS targets the development of software implementing the central service and a library and template for implementing the site-local components. Each participating site is responsible to implement the site-local components that facilitate their integration with the Fenix infrastructure. The site-local components handle aspects such as user management, resource allocation, and accounting (for further details please see ICEI Deliverable D3.3[24]).

An important focus in WP6 is the adaptation of the generic Fenix infrastructure services to the specific needs of EBRAINS, e.g. Authentication and Authorization Infrastructure (AAI) integration and FURMS. Up to now, users can login to the production OpenStack systems at ETHZ-CSCS, JUELICH-JSC, BSC and CEA using their EBRAINS account via the Fenix AAI. In addition, with an eye to additional systems and services that will come online, the EBRAINS account login via Fenix AAI (see Figure 3) has been validated on the FURMS integration environment (for further details please see Output OP6.3[25]).

[22] P1949: Leah E. Lackner, Hamid Mohammadi Fard, Felix Wolf: Efficient Job Scheduling for Clusters with Shared Tiered Storage. In Proc. of the 19th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid), Larnaca, Cyprus, pages 321–330, IEEE, May 2019.
[23] https://github.com/HumanBrainProject/coallocation-slurm-plugin
[24] This Deliverable will be made available on the Fenix website once accepted: https://fenix-ri.eu/about-fenix/deliverables
The ICEI Project Coordination Office (icei-coord@fz-juelich.de) can be contacted for a preliminary version.
[25] https://drive.ebrains.eu/f/bc54c2867b894d61a880/

Progress summary and next steps

FURMS is developed in four phases with the final phase delivering the final version of the software scheduled for early 2022. The site-local agents are currently under development and should be finalised within the first quarter of 2022.

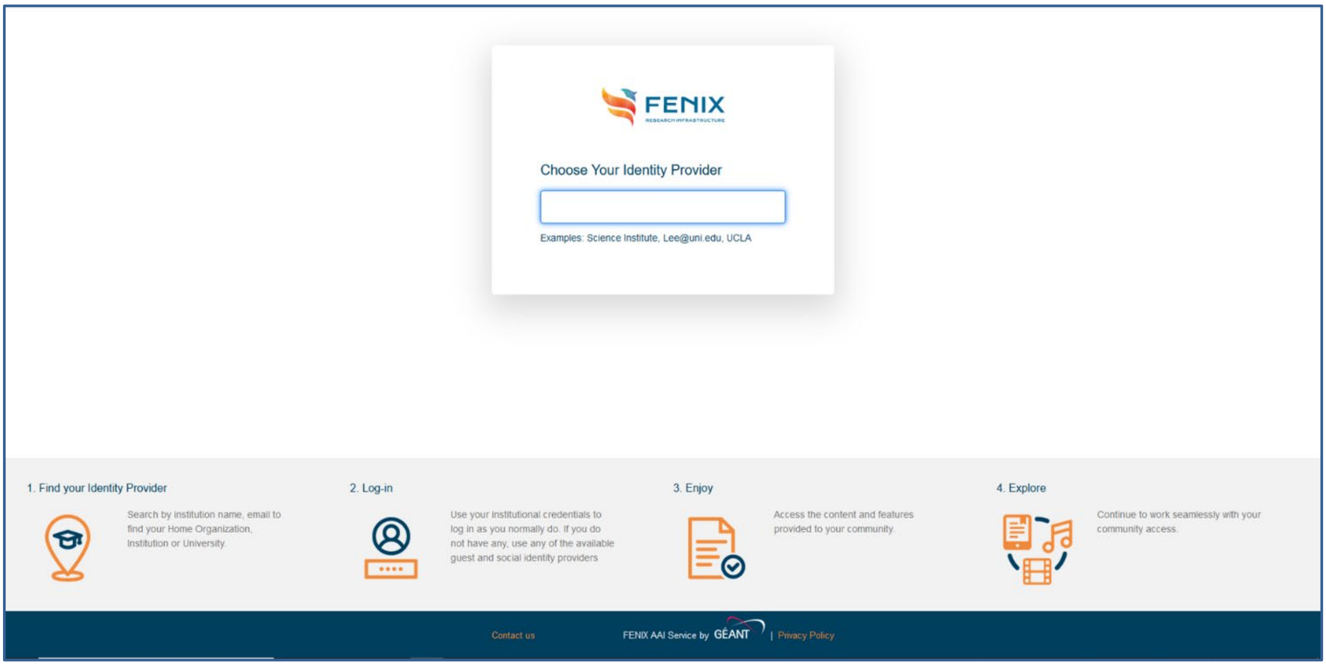

Figure 3: Fenix AAI

3.1.5 Production-level monitoring service

A production-level service for monitoring the operation of the distributed e-infrastructure services and for incident reporting was put in place based on Icinga[26] and is available[27] to every EBRAINS user. This service is of special interest for platform administrators and users who directly use infrastructure services. Using this tool, EBRAINS infrastructure users will be able to see the status and healthiness of any of the infrastructure hosts and services of the federated sites. It will also keep a log of any issue detected, which will help to debug any issue or outage, whether it was a platform-related or an infrastructure-related cause. A notification feature has been added so that interested users may get an email notification every time that a specific infrastructure service is suffering an outage or issue, and also upon its corresponding recovery.

The service is read-accessible for everyone with an EBRAINS account, while members of the related Task T6.7 have additional privileges on this website. A dashboard shows current incidents at service and host level, as well as recently recovered services. The service grid (see Figure 4) gives a good overview of the status of the services at the different hosting sites. More metrics will be added in the next weeks and months.

[26] https://icinga.com/
[27] https://wiki.ebrains.eu/bin/view/Collabs/ebrains-computing-services-wp6-public/Monitoring%20service/

Progress summary and next steps

- Monitoring service based on Icinga has been deployed, all infrastructure sites are reporting main hosts and services to this monitoring tool.
- Notification functionality has been implemented for Neuromorphic Computing hosts and services.
- List of metrics and checks for monitoring infrastructure components will be extended.
- The monitoring tool configuration will need to be updated in synchronisation with any infrastructure component change.

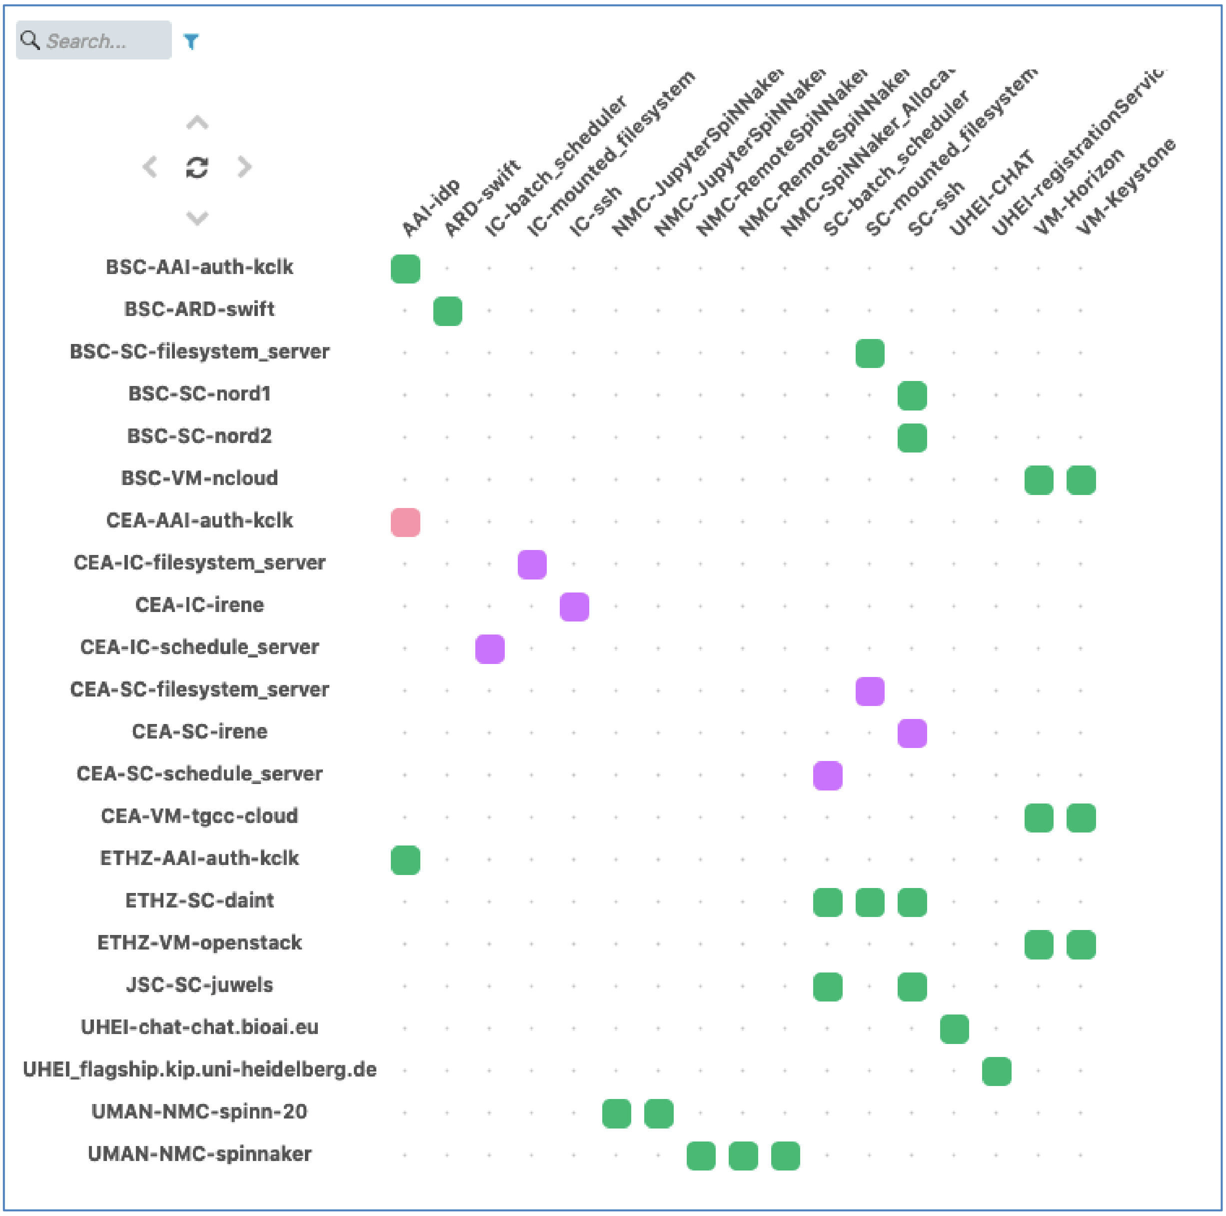

Figure 4: Service grid of the monitoring website

This service grid shows the status of all monitored services on the respective hosting systems. Green stands for “no issues / service is running”, purple means “unknown” (i.e. no information about this service is available), light pink points to a current issue. A mouse-over for every coloured box gives more information about the respective status.

3.1.6 Infrastructure usable for interactive computing workloads

An interactive computing (IAC) service is offered within EBRAINS by the five ICEI sites which also provide the underlying infrastructure as part of Fenix, namely BSC, CEA, CINECA, ETHZ-CSCS and

JUELICH-JSC. At all sites the planned procurements progressed, and the required IAC resources are available[28] now.

All deployed systems provide generic hardware and software capabilities that can handle a sufficiently wide range of usage scenarios well. All systems offer GPUs for visualisation and computing purposes. The IAC resources at BSC and CEA are particularly well suited for lowly-scalable applications requiring a large node-addressable memory space. Most IAC services are collocated with resources designed for batch-processing, which are called scalable computing services (SCC) in Fenix. The software used at the sites to interact with the deployed IAC resources is comparable, but there are relevant differences in particular with regards to the scientific software stacks.

The IAC service is designed based on an in-depth analysis of use cases gathered and described in the public ICEI deliverable D3.6 "Scientific Use Case Requirements Documentation"[29]. In contrast to previous assessments, this analysis was focused on the extraction of information for the configuration and optimisation of the deployed systems rather than the development of future components, because the SGA3 Task in charge (T6.9) of the IAC service is focused on operations.

While from a user's perspective there is a large overlap of IAC and SCC, the IAC service can be tuned separately to accommodate specific requirements of the use cases. The IAC resources at all sites use the workload manager Slurm for job scheduling and resource management. The heterogeneous job feature of Slurm[30] supports the co-scheduling of SCC and IAC jobs, so that both resources are available at the same time. Spawned processes of the job are also placed in the same process management namespace, which allows MPI applications to use intra- and inter-communicators to send messages between both types of resources.

An important goal of the interactive computing service is to speed up the resource allocation to guarantee interactivity for the user in contrast to a batch-based usage mode. The approach taken to achieve this is to analyse the batch system behaviour and to propose operation enhancements of the batch system Slurm. Aspects to work on in particular are reservations, separate job partitions as well as job boosting and pre-emption.

The IAC service also enables access to compute resources orchestrated from a web service, which itself is intended to provide an interactive experience to its users. The integration of a web service or other platforms is possible using UNICORE, based on work in earlier phases of the HBP. The IAC workload manager and a UNICORE TSI daemon are integrated via the UNICORE RESTful API, which makes CPU and GPU resources accessible from a web service. Reverse connections using https or secure shell protocols to select platform hosting servers are possible from the login partition of the IAC systems.

Visualisations are an important aspect pointed out in the use case descriptions. Since community-specific software is usually not pre-installed and can even be relatively difficult to install, it is possible to ship customised software environments in the form of container images and to execute them within the container runtime environments offered at all sites. Some sites also offer JupyterHub services, including browser-based visualisation, for easy access and a flexible and sophisticated interface to IAC and SCC.

We have considered a number of different options to enable the IAC service at all sites, including

- Workload manager (Slurm) configurations: scheduler parameters, resource dedication (at some sites IAC and SCC share the same resources, which can lead to competing job allocations);
- Priority shaping: prioritising interactive jobs over other pending requests, in particular SCC jobs;
- Alternative operational models: pre-emption (premature halting or termination of running jobs to free up resources for recently submitted interactive job requests), resource oversubscription.

While the workload manager is identical at all sites, there are many either technical or policy level boundary conditions that have an impact on the way the IAC service can be offered across them.

[28] https://fenix-ri.eu/infrastructure/resources/available-resources
[29] https://fenix-ri.eu/about-fenix/deliverables
[30] https://slurm.schedmd.com/heterogeneous_jobs.html

While, for example, job boosting is an easy means to speed up the scheduling of jobs, it makes the scheduling for other users of the systems unpredictable and is therefore not supported everywhere. Another example is pre-emption, which requires adjusting the jobs that are cancelled in favour of interactive jobs, because without taking snapshots a cancelled job could potentially lose a significant amount of computed results. Due to these boundary conditions and the varying offers of the IAC service, the use cases requesting it need to be well understood to offer the best possible implementation.

Progress summary and next steps

- The IAC was made available at all ICEI sites.
- The workload manager has been analysed via a benchmark that has also been implemented.
- Different parameters have been tested to enhance the overall interactivity.
- Different options have been discussed and tested to offer IAC at all sites, while fulfilling and taking all site-specific requirements and limitations into account.
- Possible use cases of the IAC have been identified.
- The service will be promoted to attract new users and workloads.

3.2 Recent and ongoing activities

3.2.1 Operation of HPC, cloud, storage and network services

Task T6.5 is responsible for the operation of the site-local HPC, cloud, storage and network services, which are deployed on top of the [Fenix infrastructure][31] at BSC, CEA, CINECA, ETHZ-CSCS and JUELICH-JSC and which provide the computational resources to the European neuroscience community via the HBP. T6.5 specifically takes care of the operation, maintenance and support of site-local systems providing network, scalable computing, cloud services, active and archival data repositories. Some of the systems are displayed in Figure 5. The operation of the interactive computing services is taken care of by Task T6.9, whose activities are described in Section 3.1.6. Another important activity is support for the platform-level services of the Collaboratory and other platforms.

As a first step, the team of T6.5 has identified the members of the operational teams at all sites and agreed on operational processes, aiming to enable a cooperation as lightweight as possible to reduce unnecessary overhead. A dedicated (access restricted) collab and monthly operations meetings are the central points for sharing information. A Kanban board is used to track relevant issues. At least one member of the operational teams of all sites participates in the "Technical Coordination Weeklies" to exchange operation-level information with platform providers and developers of EBRAINS services.

Already in earlier phases of the HBP, a ticketing system was put in place and operated by BSC for handling support requests related to the High Performance Analytics and Computing (HPAC) Platform, which is now part of EBRAINS Computing Services. This ticketing system is still in place and connected to the EBRAINS ticketing system, so that all tickets concerning the HPC, cloud, storage and network services are forwarded to the HPAC ticketing system and handled there. Information on the ticket status are sent back to the EBRAINS ticketing system.

More details about the operation of HPC, cloud, storage and network services can be found in a public version of the report "Operational teams defined and operational processes defined"[32] (related to Milestone MS6.1).

[31] https://fenix-ri.eu/infrastructure/resources/available-resources
[32] https://drive.ebrains.eu/f/9e3f20bd4889476a9309/?dl=1

Figure 5: ICEI systems

The picture shows some of the ICEI systems operated by BSC, CEA, CINECA, JUELICH-JSC and ETHZ-CSCS (left to right, top to bottom).

Progress summary and next steps

- In the early months of the current project phase, and as detailed in MS6.1 and in the summary above, operational members at all the sites were identified, and monthly meetings setup to facilitate exchange. Furthermore, key meetings in HBP, such as the Technical Coordination Weeklies, were identified that representatives from the sites should attend to exchange information with platforms and developers in HBP/EBRAINS.
- To support use cases such as the Brain Simulation Platform and the Brayns MOOC, the sites agreed upon a standardised checklist to enable outgoing SSH connections.
- The sites have actively responded to and resolved tickets received via the HPAC/EBRAINS ticketing system, from a wide variety of users and use cases.
- A number of key use cases have been discussed, and also how from an infrastructure level we can best support these, including but not limited to The Virtual Brain, the Neurorobotics Platform, the Medical Informatics Platform (MIP) and Human Intracerebral EEG Platform (HIP).
- In the coming months, as the infrastructure services at all sites go into production, we will continue to engage with and support HBP/EBRAINS platforms and users, as well as other activities in HBP/EBRAINS to promote infrastructure services to potential new users.

3.2.2 Security assessment of the HPC, Cloud and Neuromorphic infrastructure

An important activity at the start of the SGA3 phase was the definition of a trust zone between the involved infrastructure partners, namely all ICEI sites (BSC, CEA, CINECA, ETHZ-CSCS and JUELICH-JSC) and the neuromorphic sites UHEI (hosting the BrainScaleS system) and UMAN (hosting SpiNNaker). To achieve this, a template for security self-assessments was agreed on based on the SCI-V2 document of the WISE community[33] and on the Fenix Security Measures Catalogue. The following levels of security have been defined to rank the SCI-V2 items:

- Level 0: Function or feature not implemented
- Level 1: Function or feature exists, is operationally implemented but not documented

[33] https://wise-community.org/

- Level 1.5: Function or feature exists, is operationally implemented and some relevant documentations exist
- Level 2: Function or feature is comprehensively documented and operationally implemented
- Level 2.5: Function or feature implemented, documented, and reviewed by an independent internal body
- Level 3: Function or feature implemented, documented, and reviewed by an independent external body
- N/A: Not applicable

For a tabular summary of the actual self-assessments, see Figure 6.

Category	Min	CEA	BSC	CINECA	ETHZ-CSCS	FZJ	BrainScaleS	SpiNNaker
OS1	1.5	ok	ok	ok	ok	ok	ok	ok
OS2	1.5	ok	ok	ok	ok	ok	NA	ok
OS3	1.5	ok	ok	ok	ok	ok	NA	ok
OS4	1.5	ok	ok	ok	ok	ok	NA	ok
OS5	1.5	ok	ok	ok	ok	ok	NA	ok
OS6	1.5	ok	ok	ok	ok	ok	NA	ok
OS7	1.5	ok	ok	ok	ok	ok	NA	ok
OS8	1.5	ok	ok	ok	ok	ok	ok	ok
OS9	1.5	ok	ok	ok	ok	ok	NA	ok
OS10	1.5	ok	ok	ok	ok	ok	NA	ok
IR1	1.5	ok	ok	ok	ok	ok	NA	ok
IR2	1.5	ok	ok	ok	ok	ok	NA	ok
IR3	1.5	ok	ok	ok	ok	ok	NA	ok
IR4	1.5	ok	ok	ok	ok	ok	NA	ok
TR1	1.5	ok	ok	ok	ok	ok	NA	ok
TR2	1.5	ok	ok	ok	ok	ok	NA	ok
TR3	1.5	ok	ok	ok	ok	ok	NA	ok
PRU1	1.5	ok	ok	ok	ok	ok	ok	ok
PRU2	1.5	ok	ok	ok	ok	ok	ok	ok
PRU3	1.5	ok	ok	ok	ok	ok	ok	ok
PRC1	1.5	ok	ok	ok	ok	ok	ok	ok
PRC2	1.5	ok	ok	ok	ok	ok	ok	ok
PRC3	1.5	ok	ok	ok	ok	ok	ok	ok
PRC4	1.5	ok	ok	ok	ok	ok	NA	ok
PRC5	1.5	ok	ok	ok	ok	ok	NA	ok
PRC6	1.5	ok	ok	ok	ok	ok	NA	ok
PRS1	1.5	ok	ok	ok	ok	ok	NA	ok
DP1	1.5	ok	ok	ok	ok	ok	ok	ok
DP2	1.5	ok	ok	ok	ok	ok	NA	ok

Figure 6: SCI-V2 self-assessment summary

A field is marked “ok” for all items where sites are above the defined minimal level of security (1.5).

In the beginning of 2020, the ICEI project published the “Fenix Security Measures Catalogue” (which is part of the updated ICEI Ethics Deliverable D6.1 (V1.2) and will be soon publicly available on the Fenix RI website[34], which is another self-assessment prepared by this project. For the Fenix Resource Providers, a “Compliance Catalogue for Secure Data Storing” has been defined. Only Active Data Repositories or Archival Data Repositories that comply can be used for projects involving the processing of personal data. Therefore, the team has decided to carry out a second self-assessment by checking the compliance with these Fenix security measures and related security items. The full outcome of this self-assessment is only available to the involved partners.

[34] https://fenix-ri.eu/about-fenix/deliverables

With these two self-assessments and the security level achieved based on those, our team that is working on “Policy management and security” managed to define a security trust zone by knowing how security is handled at the other sites. To briefly summarise the outcomes, all five ICEI sites have the same goals and are equivalent. The two other sites UHEI and UMAN are hosting more specific systems dedicated to special applications, but they also completed the same self-assessments.

The self-assessments are only snapshots of the security levels of each site at a given point in time. In order to maintain the security level (and thus also the confidence in the trust zone), the team recommends and plans that each site repeat these assessments every year. In addition, work on a template to assess the security of a service has started, based on a document provided by PRACE, and also a Common Acceptable Use Policy (AUP) is work in progress.

We have set up a small Computer Security Incident Response Team (C-SIRT), which operates as a first point of contact for (potential) security issues. This official HBP C-SIRT contact will inform the other partners in case of attacks or incidents.

The team also gives advice on security-related topics for other HBP teams, e.g. to develop a security concept for handling personal data when using infrastructure resources and services. It also contributes actively to the HBP-wide “EBRAINS Infrastructure Security Working Group”.

Progress summary and next steps

- The Policy management and security team will increase and update the minimal security level required for the infrastructure. The SCI V2 assessment must be updated (as the sites' environments evolve) and completed by including different constraints found during feedbacks (in 2022).
- The Fenix Security Measures Catalogue is based on the C5 version of 2016. It will be rebased on the C5 version of 2020 (in Q1/2022).
- Having this up-to-date security reference, the infrastructure will define minimal criteria for the services running on it. To this end, a service description template will be defined (end of 2021) and completed by all services running on the infrastructure (in Q2/2022). Based on this information, minimal security criteria will be defined (in Q3/2022).
- To push the assessments forward, a process to implement self-certification will be studied (in Q3/2022).
- The Policy management and security team will continue to give advice on security-related topics (during 2022).

3.2.3 Identification of a prioritised set of selected workflows

Important aspects to ensure that the federated infrastructure provides and prioritises the services and resources that are most needed by the user community, are to be in regular contact with this community, keep gathering and analysing their evolving use cases, and select important, representative use cases to guide the developments and finally set priorities, if necessary. The EBRAINS Scientific Liaison Unit (SLU acts at the interface between developers and existing users and guides this process. Specifically, the SLU focuses on the showcases and other identified high priority science/use cases, which have high potential to be reused by other research groups.

The SLU used different sources for identifying science cases to derive workflows from them, e.g. the set of showcases defined by the scientific Work Packages of the HBP and the winning proposals of the HBP's Calls for Expression of Interest. In a first attempt to identify use cases and resulting EBRAINS workflows, the SLU has selected and prioritised the following set of workflows (see Table 2).

The SLU has collected, analysed and prioritised workflows, specifically focusing on the development of integrated workflows that are composed of multiple EBRAINS services. The first step was the definition of a template used to collect and quantify technical requirements of the infrastructure use cases, which was done in collaboration with technical and scientific experts. This template is based on a prior requirements assessment of the ICEI project (see ICEI Deliverable D3.6 “Scientific

Use Case Requirements Documentation"[35]), and an earlier version was used during SGA2. The template is filled for every use case by a team composed of the scientific use case owner and supporting technical experts, moderated by the SLU. The main steps and roles of the iterative process to describe of EBRAINS workflows in a standardised way are illustrated in Figure 7.

Table 2: Set of workflows identified by the EBRAINS Scientific Liaison Unit

Title	Priority
Degeneracy in neuroscience - when is Big Data big enough? (Showcase 1; WP1)	1
Improving epilepsy surgery with the Virtual BigBrain (Showcase 2; WP1)	1
Brain Complexity and Consciousness (Showcase 3; WP2)	1
Object Perception and Memory (Showcase 4; WP2)	1
Dextrous manipulation - how the brain coordinates hand movements (Showcase 5; WP3)	1
Calls for Expression of Interest	2
External users of the EBRAINS infrastructure as generated by Outreach (WP8), the community building tasks in the science (WP1-3) and infrastructure (WP4-6) Work Packages.	2
Existing workflows with high potential for reuse e.g. • Live papers • Brain signal analysis workflow • Full service MOOGs for EBRAINS components and science cases	2

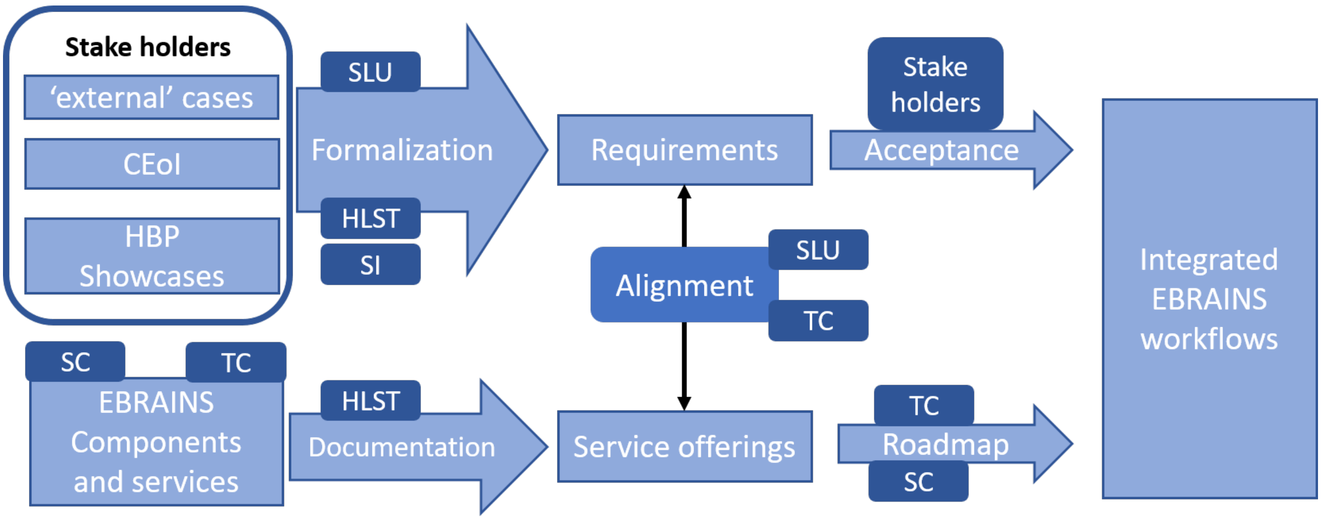

Figure 7: Steps and roles in the formalisation process of the SLU

A more detailed introduction of the SLU, its relationship with other coordination Tasks in the HBP, challenges encountered in formalising science and engineering cases, the actual template used to define the use cases, the methods and procedures followed to select the workflows, and the identified challenges are described in the report "Prioritised set of selected workflows, leading to the development of solutions based on a composition of multiple EBRAINS services" (OP6.31[36]). This report constitutes the first step in a multi-year effort. The follow-up report "First workflow solution made public to the scientific community" (OP6.32) planned for February 2022 will include the introduction of the first workflow to the scientific community. With OP6.32 the integration of the showcases into cross-service category workflow solutions begins.

Progress summary

- Formalisation process has started for the five original showcases and two additional cases, i.e. the brain signal analysis and the Learning-2-Learn workflows.

[35] https://fenix-ri.eu/about-fenix/deliverables
[36] https://drive.ebrains.eu/smart-link/09725f77-af67-427a-8b98-bc9393cf855f/

- Formalisation process has started for the Brain signal analysis workflow.
- Requirements analysis and first implementation specifically for Showcase 5 has started in collaboration with the EBRAINS Technical Coordination.
- OP6.31 compiled and delivered (see above).
- Aggregate requirements spanning multiple showcases actively placed on work backlog of EBRAINS; progress monitoring and prioritisation are managed by the SLU team.
- 58 tickets successfully closed in the SLU ticketing system.
- Outreach related activities during conferences and EBRAINS internal events.

Next steps

- Investigation and design for a cross EBRAINS workflow combining elements of all showcases.
- Further formalisation of user requirements analysis methodology.
- Requirements analysis and follow-up activities for high priority cases that are not yet addressed.

3.2.4 Infrastructure trainings and webinars

The scientists and developers in WP6 regularly participate and contribute to scientific events with seminars, workshops, presentations and posters. The following is only a selection of recent events.

We have created a public collab[37] for interested parties to find material on the services and developments of WP6, which has been created and collected for the different events. It was used for the first time for the HBP Summit 2021 and will also be used for future events.

In October 2020, our Work Package invited everyone interested in Scalable Computing resources, Virtual Machines, Neuromorphic Computing, and the Collaboratory to the first “EBRAINS Infrastructure Training Event: Tools and Services developed by EBRAINS Computing Services”. With this event an opportunity was provided for scientists, both HBP-internal and external, at different stages of their careers, to meet and learn about recent developments in and around the e-infrastructure services of EBRAINS. Public lectures provided an overview of the EBRAINS platform and insights into how one’s research could benefit from the offered e-infrastructure services. The hands-on sessions provided a guided online experience of some of the available EBRAINS services. Due to the worldwide Covid-19 pandemic this event was held fully virtually; recordings of all public sessions are online accessible[38]. In addition to the scientific topics, two of our professors gave insights into their scientific careers in a Career Coffee Break and answered the questions of the young generation. In total 95 people attended the event during the 2.5 days. Further infrastructure trainings are planned.

In November 2020, a Fenix AAI/FURMS Workshop provided an overview of the current status and developments. The workshop was also used to define functional and performance objectives and metrics for both Fenix AAI and FURMS for the execution of EBRAINS workflows. In total 33 experts attended the workshop.

Also in November 2020, the HBP CodeJam#11 took place. Here, several workshops were organised by our experts, such as a workshop on PyTorch for spiking neural networks, PyNN in data-heavy high-performance applications and Sustainable software environments for computation-heavy science, as well as a SpiNNaker-NRP Synchronisation Session, an OpenStack Tutorial, and a Scalable Compute Tutorial.

Various webinars[39] were organised together with the ICEI project. For example, the 12th Fenix webinar presented the remote access services for neuromorphic computing using the EBRAINS infrastructure. The webinars gave insights into how to use the services from a user perspective and

[37] https://wiki.ebrains.eu/bin/view/Collabs/sc6-booth
[38] https://www.humanbrainproject.eu/en/ebrains-tools-2020/
[39] https://fenix-ri.eu/media/webinars

explained how the services are built and deployed using ICEI OpenStack services. The presentation slides are online[40] available and a recording of the session is online accessible.

Under the theme "Pushing the Boundaries of Brain Research" the first Human Brain Project & EBRAINS Summit took place in October 2021. During the fully virtual event our experts answered questions from interested visitors at the Science Market. In addition, we organised a number of sessions during the internal and public days of the four-day conference:

- Keynote Talk: Neuromorphic Computing and Engineering: Past, Present, and Future
- Panel discussion: The future of brain-inspired technologies
- Breakout Session: Compute and Storage Infrastructure
- Breakout Session: Neuromorphic computing
- Panel Discussion: The Concept of EBRAINS (Services) and the Future Perspective
- Parallel Session: High Performance Computing, Cloud Computing & the Collaboratory
- Parallel Session: Neuromorphic Computing integration with HPC and NRP
- Parallel Session: Introduction to data transfer service with FTS3 and open questions

4. Looking Forward

Section 3 briefly outlines the next steps for each service and activity. The final results of all developments and activities will be described in Deliverable D6.4 "Final release of the federated HPC, Cloud and storage infrastructure for EBRAINS" by the end of the HBP. By that time, the following activities should be concluded:

- Federated user and resource management, accounting and reporting is compliant with FURMS for EBRAINS accounting and reporting workflows
- SLURM plug-in at TRL7 is tested with support for data locality in workflows and an optimised scheduling algorithm and available at ICEI
- Workshop and training event for knowledge exchange and transfer for Fenix AAI and FURMS was organised
- Evaluation of all ICEI services
- Integration of composable EBRAINS services into cross-SC workflow solutions

WP6 will also investigate how to maintain the developed EBRAINS Computing Services and keep them available for the neuroscience community beyond the end of the current HBP SGA3 project phase.

[40] https://fenix-ri.eu/sites/default/files/public/file-uploads/fenix_webinar_nmc.pdf